

(weak) Calibration is Computationally Hard

Elad Hazan

Technion - Israel Institute of Technology
ehazan@ie.technion.ac.il

Sham M. Kakade

Microsoft Research, New England & Wharton, University of Pennsylvania
skakade@microsoft.com

Abstract

We show that the existence of a computationally efficient calibration algorithm, with a low weak calibration rate, would imply the existence of an efficient algorithm for computing approximate Nash equilibria — thus implying the unlikely conclusion that every problem in *PPAD* is solvable in polynomial time.

1 Introduction

Consider a weather forecaster that predicts the probability of rain. The forecaster is said to be *calibrated* if every time she predicts a certain probability of rain, the empirical average of rainy vs. non-rainy days approaches this forecasted probability.

This very natural property of forecasting was introduced by [Daw82] and has found numerous applications since [FV97, FV98, KLS99, Fos99, FL99, MSA07, Per09, MS10, RST11]. See [CL06] for a more detailed bibliographic survey.

[FV98] provided the first randomized calibration algorithms. Subsequently, numerous other algorithms have been developed based on various different techniques have followed: Blackwell approachability [Fos99], internal-regret minimization [FV98] and online convex optimization [ABH11], to name a few.

While existence results for calibration are well established, our understanding of the statistical and computational complexity is more murky. The statistical complexity can be thought of as the number of rounds it takes achieve some natural notion of a low calibration; the computational complexity can be thought of as the net computation time to achieve this. This work provides a lower bound for the latter. When characterizing the efficiency of algorithms, the critical issue is the relationship between the relevant parameters and the desired notion of calibration. The notion of the (total) calibration rate (at precision ε) is essentially that defined by [FV98]. The relevant parameters are the number of forecasting iterations (henceforth denoted T), the precision of calibration ε , and number of possible outcomes in the forecasting game, d . A variant of this question was posed as an open problem in [AM11].¹

¹[AM11] did not explicitly pose this question in terms of net computation time.

In this work, we give a negative result showing that calibration (in the worst case) is hard, under a widely-believed computational complexity assumption. In particular, we utilize a natural (smooth) notion of calibration at scale ε , namely *weak calibration* (as in [KF08]). Precisely, the complexity implication of our main result, Theorem 3, is as follows:

Corollary 1. *Suppose there exists a constant $c > 0$ and a weak calibration algorithm which, for every precision $\varepsilon > 0$, attains a calibration rate of ε^c in a total computational running time (in the RAM model) that is polynomial in both d and $\frac{1}{\varepsilon}$, then $PPAD \subseteq RP$.*

Here, the weak calibration rate is a cumulative notion of error, precisely defined in Section 2; RP stands for the complexity class of randomized polynomial time; $PPAD$ is the class of problems that are polynomial time reducible to the problem of computing Nash equilibrium in a two player game (See [Pap94, Das09]). It is widely believed that $PPAD$ is not contained in RP . Note that we are considering the *total* computation time over all T rounds (so there is no explicit T dependence).

2 Calibration

Calibration inherently concerns distributions, and when comparing distributions it makes sense to talk about statistical distance or its closely related cousin the ℓ_1 norm, rather than the Euclidean norm. Therefore throughout we use $\|\cdot\|$ to denote the ℓ_1 norm and $\|\cdot\|_p$ to denote the ℓ_p norm.

We let $\{0, 1, 2, \dots, d\}$ be an outcome space, and X_1, X_2, \dots, X_T be a sequence of outcomes, denoted as $X_t \in \{0, 1\}^d$, such that $X_t(i)$ is one if and only if the outcome in iteration t is $i \in [d]$. Hence $\frac{1}{T} \sum_t X_t$ is the empirical frequency of outcomes.

A randomized forecaster \mathcal{A} produces a sequence of probability distributions $\mathcal{D}_1, \dots, \mathcal{D}_T$ over the set $\Delta_d = \{p \in \mathbb{R}^d, p_i \geq 0, \sum_i p_i = 1\}$. Every iteration a point in the interior of the simplex is chosen: $p_t \sim \mathcal{D}_t$, which constitutes the forecast of \mathcal{A} .

Strong Calibration: For a set of points $V \subset \Delta_d$, define the following “test” functions (where the arg min breaks ties arbitrarily):

$$\mathbb{I}_p(q) = \begin{cases} 1 & p = \arg \min_{p' \in V} \|p' - q\| \\ 0 & \text{otherwise} \end{cases}$$

We say this set of test function is at *precision* ε if V is such that every $q \in \Delta_d$ is at least ε -close (in ℓ_1) to some point in V , i.e. for all $q \in \Delta_d$, we have $\min_{p \in V} \|p - q\| \leq \varepsilon$ (i.e. the set V is an ε -cover for Δ_d).

Definition 1. *Let the strong-calibration rate of a (possibly randomized) forecaster \mathcal{A} , with respect to indicator test functions $\mathcal{F}^\varepsilon = \{\mathbb{I}_q(\cdot)\}$ at precision ε , be*

$$C_T(X_{1:T}, \mathcal{A}, \mathcal{F}^\varepsilon) = \mathbb{E}_{\mathcal{D}_1, \dots, \mathcal{D}_T} \left[\frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{I}_p(p_t) (p_t - X_t) \right\| \right]$$

This definition is closely related to that used in [BL85, FV98]; the latter definition is motivated by a bias-variance decomposition of the Brier score. The distinctions being that [FV98] use the squared ℓ_2 error (while we use the ℓ_1 primarily for convenience) and [FV98] restrict \mathcal{A} to make predictions which lie in V (a minor distinction).

Much of the literature is concerned with the asymptotic behavior, without explicitly characterizing the finite time rate. It is standard to say that a forecaster \mathcal{A} is (strongly) *asymptotically calibrated* if for all $X_{1:T}$, we can drive $C_T(\mathcal{A}, \mathcal{F}^\varepsilon)$ to 0, as $T \rightarrow \infty$. If \mathcal{A} is restricted to make predictions in the set V , then this notion seeks to drive $C_T(\mathcal{A}, \mathcal{F}^\varepsilon) \leq \varepsilon$ in the limit. In this work, the rate of this function is critical.

The definition of asymptotic calibration considers the “total error” over an ε -grid, and it adjusts the normalization for each term to $\frac{1}{T}$. Note that our indicator functions satisfy for all $q \in \Delta_d$:

$$\sum_{p \in V} \mathbb{I}_p(q) = 1 \quad (1)$$

Since every q is covered by only one indicator function. This implies that:

$$\frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \mathbb{I}_p(p_t) = 1$$

which implies that $C_T(X_{1:T}, \mathcal{A}, \mathcal{F}^\varepsilon)$ is bounded by 2.

Weak Calibration: We now turn to the notion of weak calibration, which covers Δ_d in a more continuous manner. The weak calibration rate is more naturally defined by a triangulation of the simplex, Δ_d . By this, we mean that Δ_d is partitioned into a set of simplices such that any two simplices intersect in either a common face, common vertex, or not at all. Let V be the vertex set of this triangulation. Note that any point q lies in some simplex in this triangulation, and, slightly abusing notation, let $V(q)$ be the set of corners for this simplex. Note that the function $V(\cdot)$ specifies the triangulation.

Instead of indicator functions $\mathbb{I}_p(\cdot)$, we associate a test function $\omega_p(\cdot)$ with each $p \in V$ as follows. Each $q \in \Delta_d$ can be uniquely written as a weighted average of its neighboring vertices, $V(q)$. For $p \in V(q)$, let us define the test functions $\omega_p(q)$ to be these linear weights, so they are uniquely defined by the linear equation:

$$q = \sum_{p \in V(q)} \omega_p(q) p$$

For $p \notin V(q)$, we let $\omega_p(q) = 0$. We refer to this set of functions as the *triangulated test functions* with regards to $V(\cdot)$ and say that this is at *precision* ε if the diameter of the set of points $V(q)$ is less than ε for all q .

A useful property is that for all $q \in \Delta_d$,

$$\sum_{p \in V} \omega_p(q) = 1 \quad (2)$$

since q lies in the convex hull of $V(q)$. In comparison to Equation (1), these test functions cover Δ_d in a more smooth manner: they again sum to 1, and each $\omega_p(q)$ is a continuous function (as opposed to the discontinuous indicator functions).

We now define deterministic calibration algorithms, so called “weak calibration” with regards to these Lipchitz test functions.

Definition 2. Let $\mathcal{W}^\varepsilon = \{\omega_p\}$ be a set of triangulated test functions at precision ε . The weak-calibration rate for a (deterministic) forecaster \mathcal{A} with respect to \mathcal{W}^ε

$$C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon) = \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \omega_p(p_t)(p_t - X_t) \right\|$$

[KF08] showed that there exist deterministic calibration algorithms (also see [MSA07]). Again, note the normalization property:

$$\frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \omega_p(p_t) = 1$$

which implies that $C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon)$ is bounded by 2.

3 Main Result

Our main result is based on using a calibration algorithm to compute a Nash equilibrium of a two player game. Before we state our main result, let us review the definition of an approximate Nash equilibrium, along with the attendant computational complexity results.

3.1 Nash equilibria in games

A (square) *two-player bi-matrix game* is defined by two payoff matrices $U_1, U_2 \in \mathbb{R}^{n \times n}$, such that if the row and column players choose pure strategies $i, j \in [n]$, respectively, the payoff to the row and column players are $U_1(i, j)$ and $U_2(i, j)$, respectively.

A *mixed strategy* for a player is a distribution over pure strategies (i.e. rows/columns), and for brevity we may refer to it simply as a strategy. An ε -*approximate Nash equilibrium* is a pair of mixed strategies (p, q) such that

$$\begin{aligned} \forall i \in [n], \quad p^\top U_1 q &\geq e_i^\top U_1 q - \varepsilon, \\ \forall j \in [n], \quad p^\top U_2 q &\geq p^\top U_2 e_j - \varepsilon. \end{aligned}$$

Here and throughout, e_i is the i -th standard basis vector, i.e. 1 in i -th coordinate, and 0 in all other coordinates. If $\varepsilon = 0$, the strategy pair is called a *Nash equilibrium* (NE).

For notational convenience, we slightly abuse notation by denoting the payoffs of mixed strategies as:

$$U_1(p, q) = p^\top U_1 q, \quad U_2(p, q) = p^\top U_2 q$$

The definition immediately implies that the pair (x, y) is an ε -equilibrium if and only if for all mixed strategies \tilde{x}, \tilde{y} ,

$$\begin{aligned} U_1(x, y) &\geq U_1(\tilde{x}, y) - \varepsilon, \\ U_2(x, y) &\geq U_2(x, \tilde{y}) - \varepsilon. \end{aligned}$$

Algorithm 1 Approximate NE computation via calibration algorithm \mathcal{A}

Input: calibration algorithm \mathcal{A} along with \mathcal{W}^ε on the outcome space $\{0, 1\}^d \times \{0, 1\}^d$; two player game U_1, U_2 over $\Delta_d \times \Delta_d$.

Initialize Set $\delta = \varepsilon^{1/3}$ and p_1 to be $\mathcal{A}(\emptyset)$

for $t = 1, 2, \dots, T$ **do**

Let $[p_t]_1$ and $[p_t]_2$ denote the marginal distributions of p_t with respect to the first and second coordinates (respectively).

Sample the outcome $X_t \in \{0, 1\}^d \times \{0, 1\}^d$ according to the product distribution:

$$X_t \sim \mathbf{BR}_{1,\delta}([p_t]_2) \times \mathbf{BR}_{2,\delta}([p_t]_1)$$

where $\mathbf{BR}_{i,\delta}$ is a smooth best-response function, defined in Section 4.1.

Update $p_{t+1} \leftarrow \mathcal{A}(X_1, \dots, X_t)$

end for

Sample t uniformly from $\{1, \dots, T\}$

Sample $p \in V(p_t)$ under the law $\Pr(p|p_t) = \omega_p(p_t)$.

return $\mathbf{BR}_\delta(p) = (\mathbf{BR}_{1,\delta}([p]_2), \mathbf{BR}_{2,\delta}([p]_1))$

As we are concerned with an additive notion of approximation, we assume that the entries of the matrices are in the range $[0, 1]$. In particular this implies that the functions U_1, U_2 are 1-Lipschitz w.r.t the ℓ_1 norm, since for all $p_1, p_2, q \in \Delta_d$:

$$U_i(p_1, q) - U_i(p_2, q) = (p_1 - p_2)^\top U_i q \leq \|p_1 - p_2\| \|U_i q\|_\infty \leq \|p_1 - p_2\| \quad (3)$$

Where we used Hölder's inequality and the fact that $U_i(i, j) \in [0, 1]$.

The following theorem was provided by [CDT09]:

Theorem 2. [CDT09] *If there exists a randomized algorithm that computes a ε -NE in a two player game in time $\text{poly}(d, \frac{1}{\varepsilon})$ then $\text{PPAD} \subseteq \text{RP}$.*

3.2 Nash equilibria computation with a calibration algorithm

We now present the reduction from weak calibration to computing equilibria in games, thereby obtaining the hardness result stated in Corollary 1. Algorithm 1 utilizes a calibration algorithm in a specially tailored game theoretic protocol. Observe this protocol is run with an outcome space of size d^2 . This protocol is based on the ideas in [KF08], which utilized a weak calibration algorithm to obtain asymptotic convergence to the convex hull of Nash equilibria (also see [MSA07]). Here, our algorithm outputs a particular approximate Nash equilibrium in finite time, which allows us to provide a computational complexity lower bound.

Theorem 3. *Suppose a weak calibration algorithm \mathcal{A} satisfies the following uniform bound on the calibration rate: $C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon) \leq F(d, \mathcal{W}^\varepsilon, T)$ (where F does not depend on $X_{1:T}$). Let $d > 2$ and $\varepsilon < \frac{1}{d^3}$. Then with probability greater than $1/2$, Algorithm 1 (using $\delta = \varepsilon^{1/3}$) returns a $(4F(d^2, \mathcal{W}^\varepsilon, T) + 22d\varepsilon^{1/3})$ -Nash equilibrium.*

This directly implies Corollary 1 as follows:

Corollary 1. Let \mathcal{A} be a weak calibration algorithm that attains a calibration rate of ε^c at precision ε . Then for some T (where T is polynomial in ε, d) we have that $C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon) \leq F(d^2, \mathcal{W}^\varepsilon, T) \leq \varepsilon^c$. Theorem 3 implies that Algorithm 1 returns a $O(\varepsilon^c + d\varepsilon^{1/3})$ -NE after T iterations with probability greater than $\frac{1}{2}$. This constitutes a randomized polynomial time algorithm for ε -NE, which by Theorem 2 implies $PPAD \subseteq RP$. \square

4 Analysis

Our analysis is arranged into three parts. First, we define a smooth best response function \mathbf{BR}_δ along with some technical lemmas. Then we show how fixed points of this \mathbf{BR}_δ function are approximate Nash equilibria. With these lemmas, we complete the proof.

4.1 Smooth Best Response Functions

Our algorithm utilizes smooth best response functions. For a mixed strategy $q \in \Delta_d$, define the best response functions as:

$$\mathbf{BR}_i(q) = \arg \max_{p \in \Delta_d} \{U_i(p, q)\}$$

In case the RHS is a set, define \mathbf{BR}_i as an arbitrary member of the set.

We say that a function $g : \Delta_d \mapsto \Delta_d$ is an ε -best response with respect to U_i if the following holds:

$$\forall q, U_i(g(q), q) \geq U_i(\mathbf{BR}_i(q), q) - \varepsilon$$

It is be convenient to extend the best response function beyond the simplex. Define for any point in Euclidean space:

$$\forall p \in \mathbb{R}^n. \mathbf{BR}_i(p) = \mathbf{BR}_i(\prod_{\Delta_d}(p))$$

where $\prod_{\mathcal{K}}(p)$ denotes the projection operation onto a convex set \mathcal{K} defined as:

$$\prod_{\mathcal{K}}(p) = \arg \min_{q \in \mathcal{K}} \|p - q\|_2$$

Using the generalized definition of \mathbf{BR}_i , define the δ -smooth best response function as:

$$\mathbf{BR}_{i,\delta}(q) := \mathbb{E}_{\|q' - q\|_\infty \leq \delta} [\mathbf{BR}_i(q')] \quad (4)$$

where the expectation is with respect to the random q' sampled uniformly on the set $\{q' \mid \|q' - q\|_\infty \leq \delta\}$.

Lemma 4. The function $\mathbf{BR}_{i,\delta}$ is a $(2d\delta)$ -best response with respect to U_i .

Proof. Let q, q' be such that $\|q - q'\|_\infty \leq \delta$. Hence, $\|q' - q\| \leq d\delta$ and since U_i is 1-Lipschitz with respect to the ℓ_1 norm (see equation (3)):

$$\forall p. |U_i(p, q') - U_i(p, q)| \leq \|q' - q\| \leq d\delta$$

Let $q' = \arg \min_{\tilde{q} \in \Delta_d, \|\tilde{q} - q\|_\infty \leq \delta} U_i(\mathbf{BR}_i(\tilde{q}), q)$. Using the definitions above, we have

$$\begin{aligned} U_i(\mathbf{BR}_{i,\delta}(q), q) &= U_i \left(\mathbb{E}_{\|q' - q\|_\infty \leq \delta} [\mathbf{BR}_i(q')], q \right) \\ &\geq U_i(\mathbf{BR}_i(q'), q) \\ &\geq U_i(\mathbf{BR}_i(q'), q') - d\delta && \text{since } \|q' - q\|_\infty \leq \delta \\ &\geq U_i(\mathbf{BR}_i(q), q') - d\delta && \text{definition of } \mathbf{BR}_i \\ &\geq U_i(\mathbf{BR}_i(q), q) - 2d\delta && \text{since } \|q' - q\|_\infty \leq \delta \end{aligned}$$

which completes the proof. \square

Lemma 5. For $2 < d < \frac{1}{\delta}$, the function $\mathbf{BR}_{i,\delta}$ is $\frac{2}{\delta^2}$ -Lipschitz.

Proof. Consider any two distributions p, q . We consider two cases:

case 1: $\|p - q\|_\infty > \delta^2$. In this case we have

$$\begin{aligned} \|\mathbf{BR}_{i,\delta}(p) - \mathbf{BR}_{i,\delta}(q)\| &\leq \|\mathbf{BR}_{i,\delta}(p)\| + \|\mathbf{BR}_{i,\delta}(q)\| && \text{triangle inequality} \\ &\leq 2 && \text{the range of } \mathbf{BR}_{i,\delta} \text{ is } \Delta_d \\ &\leq \|p - q\|_\infty \cdot \frac{2}{\delta^2} && \text{by condition on } \|p - q\|_\infty \\ &\leq \|p - q\| \cdot \frac{2}{\delta^2} \end{aligned}$$

case 2: $\|p - q\|_\infty \leq \delta^2$. Denote the d -dimensional cube with radius δ centered at p by

$$\mathcal{C}_\delta^d(p) = \mathcal{C}_\delta(p) = \{q \in \Delta_d, \|q - p\|_\infty \leq \delta\}$$

We have

$$\begin{aligned} \|\mathbf{BR}_{i,\delta}(p) - \mathbf{BR}_{i,\delta}(q)\| &= \left\| \mathbb{E}_{\|p' - p\|_\infty \leq \delta} [\mathbf{BR}_i(p')] - \mathbb{E}_{\|q' - q\|_\infty \leq \delta} [\mathbf{BR}_i(q')] \right\| \\ &= \left\| \mathbb{E}_{p' \in \mathcal{C}_\delta(p)} [\mathbf{BR}_i(p')] - \mathbb{E}_{q' \in \mathcal{C}_\delta(q)} [\mathbf{BR}_i(q')] \right\| \\ &\leq \frac{\text{vol}(\mathcal{C}_\delta(p) \setminus \mathcal{C}_\delta(q) \cup \mathcal{C}_\delta(q) \setminus \mathcal{C}_\delta(p))}{\text{vol}(\mathcal{C}_\delta(p) \cup \mathcal{C}_\delta(q))} \\ &\leq 2 \frac{\text{vol}\{\mathcal{C}_\delta(p) \setminus \mathcal{C}_\delta(q)\}}{\text{vol}(\mathcal{C}_\delta(q))} \end{aligned}$$

The volume of $\mathcal{C}_\delta(x)$ for any $x \in \mathbb{R}^d$ is given by δ^d . To bound the volume of $\mathcal{C}_\delta(p) \setminus \mathcal{C}_\delta(q)$ notice that at least one coordinate of any point in this set is within distance δ of

p but not of q . Hence, the range of possible values for this coordinate is bounded by $\|p - q\|_\infty$. This is possible for all d coordinates, and we obtain:

$$\text{vol}\{\mathcal{C}_\delta(p) \setminus \mathcal{C}_\delta(q)\} \leq \|p - q\|_\infty \cdot d \cdot \text{vol}(\mathcal{C}_\delta^{d-1}(p)) \leq d\|p - q\|_\infty \delta^{d-1}$$

We conclude that:

$$\begin{aligned} \|\mathbf{BR}_{i,\delta}(p) - \mathbf{BR}_{i,\delta}(q)\| &\leq 2 \frac{\text{vol}\{\mathcal{C}_\delta(p) \setminus \mathcal{C}_\delta(q)\}}{\text{vol}(\mathcal{C}_\delta(q))} \\ &\leq \frac{2\|p - q\|_\infty d \delta^{d-1}}{\delta^d} \leq \frac{2d}{\delta} \cdot \|p - q\|_\infty \leq \frac{2}{\delta^2} \|p - q\|_\infty \end{aligned}$$

which completes the proof. \square

4.2 Approximate Nash equilibria and fixed points

Lemma 6. (*Approximate NE are Approximate Fixed Points*) Let p be a (possibly joint) distribution on the space of outcomes $\{0, 1\}^d \times \{0, 1\}^d$; let $[p]_1$ and $[p]_2$ denote the marginal distributions of p with respect to the first and second coordinates (respectively); let $\mathbf{BR}_\delta(p)$ denote the product distribution $\mathbf{BR}_{1,\delta}([p]_2) \times \mathbf{BR}_{2,\delta}([p]_1)$. Suppose

$$\|p - \mathbf{BR}_\delta(p)\| \leq \gamma$$

Then $\mathbf{BR}_\delta(p)$ is a $(2\gamma + 2d\delta)$ -NE.

Proof. By construction, $\mathbf{BR}_\delta(p)$ is a product distribution. Hence, it suffices to show that $\mathbf{BR}_{1,\delta}([p]_2)$ is an $(2\gamma + 2d\delta)$ -best response to $\mathbf{BR}_{2,\delta}([p]_1)$ (and vice versa). First, observe that:

$$\|[q]_1 - [p]_1\| = \sum_{i=1}^d \left\| \sum_{j=1}^d (q(i, j) - p(i, j)) \right\| \leq \sum_{i,j=1}^d \|q(i, j) - p(i, j)\| = \|q - p\| \quad (5)$$

Similarly, $\|[q]_2 - [p]_2\| \leq \|q - p\|$ Hence,

$$\|[p]_i - \mathbf{BR}_{i,\delta}(p)\| \leq \|p - \mathbf{BR}_\delta(p)\| \leq \gamma$$

By Lemma 4, $\mathbf{BR}_{1,\delta}([p]_2)$ is a $2d\delta$ -best response to $[p]_2$. Since $\|[p]_2 - \mathbf{BR}_{2,\delta}([p]_1)\| \leq \gamma$, we have that for all $q \in \Delta_d$,

$$|U_1(q, [p]_2) - U_1(q, \mathbf{BR}_{2,\delta}([p]_1))| \leq \gamma$$

Hence, for all $q \in \Delta_d$,

$$\begin{aligned} U_1(\mathbf{BR}_{1,\delta}([p]_2), \mathbf{BR}_{2,\delta}([p]_1)) &\geq U_1(\mathbf{BR}_{1,\delta}([p]_2), [p]_2) - \gamma \\ &\geq U_1(q, [p]_2) - \gamma - 2d\delta \\ &\geq U_1(q, \mathbf{BR}_{2,\delta}([p]_1)) - 2\gamma - 2d\delta \end{aligned}$$

which proves the claim. \square

5 Proof (of Theorem 3))

Three observations are helpful for intuition in the proof:

- By construction in Algorithm 1, in expectation, the outcomes X_t are just $\mathbf{BR}_\delta(p_t)$. Precisely, $E[X_t | X_1, \dots, X_{t-1}] = \mathbf{BR}_\delta(p_t)$.
- Suppose $\omega_p(p_t)$ is nonzero (so $\|p - p_t\| \leq \varepsilon$). Then, by Lemma 5, the larger δ is the closer $\mathbf{BR}_\delta(p_t)$ and $\mathbf{BR}_\delta(p)$ will be to each other.
- The smaller δ is, the more accurate an approximate NE we have for an approximate fixed point of \mathbf{BR}_δ (by Lemma 6).

The proof of Theorem 3 is a consequence from the following lemma.

Lemma 7. *Let p and $X_{1:T}$ be the random variables defined in Algorithm 1. For $2 < d < \frac{1}{\delta}$, we have that:*

$$\mathbb{E} \|p - \mathbf{BR}_\delta(p)\| \leq \mathbb{E}[C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon)] + \varepsilon + \frac{4\varepsilon}{\delta^2}$$

The proof of our Main result now follows:

Theorem 3. By Markov's inequality, we have that with probability greater than $1/2$

$$\begin{aligned} \|p - \mathbf{BR}_\delta(p)\| &\leq 2 \mathbb{E}[C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon)] + 2\varepsilon + \frac{8\varepsilon}{\delta^2} \\ &\leq 2F(d^2, \mathcal{W}^\varepsilon, T) + 10\varepsilon^{1/3} \end{aligned}$$

using the definition of F (on a d^2 sized outcome space) and $\delta = \varepsilon^{1/3}$. By applying Lemma 6, we have a $(4F(d^2, \mathcal{W}^\varepsilon, T) + 20\varepsilon^{1/3} + 2d\varepsilon^{1/3})$ -NE, which completes the proof. \square

We continue to prove Lemma 7:

Lemma 7. We proceed by lower bounding the expected calibration rate as follows:

$$\begin{aligned}
& \mathbb{E}[C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon)] \\
&= \mathbb{E} \left[\sum_{p \in V} \left\| \frac{1}{T} \sum_{t=1}^T \omega_p(p_t)(p_t - X_t) \right\| \right] \\
&\geq \frac{1}{T} \sum_{p \in V} \left\| \mathbb{E} \left[\sum_{t=1}^T \omega_p(p_t)(p_t - X_t) \right] \right\| && \text{Jensen's} \\
&= \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p_t - X_t)] \right\| && \text{linearity} \\
&= \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\mathbb{E}[\omega_p(p_t)(p_t - X_t) | X_1, \dots, X_{t-1}]] \right\| \\
&= \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p_t - \mathbf{BR}_\delta(p_t))] \right\| && p_t \text{ is determined by the history}
\end{aligned}$$

Note that by construction in Algorithm 1 $\mathbb{E}[X_t | X_1, \dots, X_{t-1}] = \mathbf{BR}_\delta(p_t)$, which we have used in the last step.

Hence, we have:

$$\begin{aligned}
\mathbb{E}[C_T(X_{1:T}, \mathcal{A}, \mathcal{W}^\varepsilon)] &\geq \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p - \mathbf{BR}_\delta(p))] \right\| \\
&\quad - \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p - p_t + \mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p))] \right\|
\end{aligned}$$

by the triangle inequality.

For the first term,

$$\begin{aligned}
& \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p - \mathbf{BR}_\delta(p))] \right\| \\
&= \frac{1}{T} \sum_{p \in V} \left\| \left(\sum_{t=1}^T \mathbb{E} [\omega_p(p_t)] \right) (p - \mathbf{BR}_\delta(p)) \right\| \\
&= \frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)] \|p - \mathbf{BR}_\delta(p)\| \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{p \in V} \omega_p(p_t) \|p - \mathbf{BR}_\delta(p)\| \right] \\
&:= \mathbb{E}_{p \sim D} \|p - \mathbf{BR}_\delta(p)\|
\end{aligned}$$

where $p \sim D$ is sampled as follows: first, sample t uniformly from $[T]$, then sample p_t according to the underlying process, and then sample $p \in V(p_t)$ with probability $\omega_p(p_t)$. Note that D is precisely the sampling procedure defined in Algorithm 1.

For the last term, we have that:

$$\begin{aligned}
& \frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p - p_t + \mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p))] \right\| \\
&\leq \frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \|\mathbb{E} [\omega_p(p_t)(p - p_t + \mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p))]\| && \text{triangle inequality} \\
&\leq \frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \mathbb{E} [\|\omega_p(p_t)(p - p_t + \mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p))\|] && \text{Jensen's} \\
&\leq \frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \mathbb{E} [\omega_p(p_t) \|p - p_t\| + \omega_p(p_t) \|\mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p)\|] && \text{sublinearity}
\end{aligned}$$

Now observe that for product distributions $D = p(x)q(y)$ and $D' = p'(x)q'(y)$.

$$\begin{aligned}
\|D - D'\| &= \sum_{x,y} |p(x)q(y) - p'(x)q'(y)| \\
&\leq \sum_{x,y} |p(x)q(y) - p(x)q'(y)| + \sum_{x,y} |p(x)q'(y) - p'(x)q'(y)| \\
&= \sum_{x,y} p(x)|q(y) - q'(y)| + \sum_{x,y} q'(y)|p(x) - p'(x)| \\
&= \|q - q'\| + \|p - p'\|
\end{aligned}$$

Also note that $V(q)$ has diameter ε , then if $w_p(q) \neq 0$ then $\|p - q\| \leq \varepsilon$. Hence,

$$\begin{aligned}
&\|\mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p)\| \\
&\leq \|\mathbf{BR}_{1,\delta}([p_t]_2) - \mathbf{BR}_{1,\delta}([p]_2)\| + \|\mathbf{BR}_{2,\delta}([p_t]_1) - \mathbf{BR}_{2,\delta}([p]_1)\| \\
&\leq \frac{2\|[p_t]_2 - [p]_2\|}{\delta^2} + \frac{2\|[p_t]_1 - [p]_1\|}{\delta^2} \quad \text{by Lemma 5} \\
&\leq \frac{4\|p_t - p\|}{\delta^2} \quad \text{by Equation 5} \\
&\leq \frac{4\varepsilon}{\delta^2}
\end{aligned}$$

where we have used Lemma 5 with our condition on d .

Hence, for the last term,

$$\begin{aligned}
&\frac{1}{T} \sum_{p \in V} \left\| \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)(p - p_t + \mathbf{BR}_\delta(p_t) - \mathbf{BR}_\delta(p))] \right\| \\
&\leq \frac{1}{T} \sum_{p \in V} \sum_{t=1}^T \mathbb{E} [\omega_p(p_t)] \left(\varepsilon + \frac{4\varepsilon}{\delta^2} \right) \\
&= \frac{1}{T} \sum_{t=1}^T \mathbb{E} \left[\sum_{p \in V} \omega_p(p_t) \right] \left(\varepsilon + \frac{4\varepsilon}{\delta^2} \right) \\
&= \varepsilon + \frac{4\varepsilon}{\delta^2}
\end{aligned}$$

The claim now follows. \square

6 Discussion and Open Problems

This work provides a computational lower bound for weak calibration, suggesting that the hardness of the problem may be fundamentally related to the problem of finding a fixed point. The following questions remain open:

- Is it possible to obtain an efficient algorithm for strong calibration? (One which gives a low calibration error in time polynomial in the relevant parameters.)
- What is the statistical complexity of (weak or strong) calibration? Here, the statistical complexity is the number of rounds required to calibrate at some desired level of accuracy, without computational considerations.

References

- [ABH11] Jacob Abernethy, Peter L. Bartlett, and Elad Hazan. Blackwell approachability and no-regret learning are equivalent. *Journal of Machine Learning Research - Proceedings Track*, 19:27–46, 2011.
- [AM11] Jacob Abernethy and Shie Mannor. Does an efficient calibrated forecasting strategy exist? *Journal of Machine Learning Research - Proceedings Track*, 19:809–812, 2011.
- [BL85] Gail Blattenberger and Frank Lad. Separating the brier score into calibration and refinement components: A graphical exposition. *The American Statistician*, 39:26–32, 1985.
- [CDT09] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player nash equilibria. *J. ACM*, 56(3):1–57, 2009.
- [CL06] Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.
- [Das09] Constantinos Daskalakis. Nash equilibria: Complexity, symmetries, and approximation. *Computer Science Review*, 3(2):87–100, 2009.
- [Daw82] A. Dawid. The well-calibrated Bayesian. *Journal of the American Statistical Association*, 77:605–613, 1982.
- [FL99] Drew Fudenberg and David K. Levine. An easier way to calibrate. *Games and Economic Behavior*, 29(1-2):131–137, October 1999.
- [Fos99] D. P Foster. A proof of calibration via blackwell’s approachability theorem. *Games and Economic Behavior*, 29(1-2):7378, 1999.
- [FV97] Dean P. Foster and Rakesh V. Vohra. Calibrated learning and correlated equilibrium. *Games and Economic Behavior*, 21(1-2):40–55, October 1997.
- [FV98] D. P Foster and R. V Vohra. Asymptotic calibration. *Biometrika*, 85(2):379, 1998.
- [KF08] Sham M. Kakade and Dean P. Foster. Deterministic calibration and nash equilibrium. *J. Comput. Syst. Sci.*, 74(1):115–130, 2008.

- [KLS99] Ehud Kalai, Ehud Lehrer, and Rann Smorodinsky. Calibrated forecasting and merging. *Games and Economic Behavior*, 29(1-2):151–169, October 1999.
- [MS10] Shie Mannor and Gilles Stoltz. A geometric proof of calibration. *Math. Oper. Res.*, 35(4):721–727, 2010.
- [MSA07] S. Mannor, J.S. Shamma, and G. Arslan. Online calibrated forecasts: Memory efficiency versus universality for learning in games. *Machine Learning*, 67(1):77–115, 2007.
- [Pap94] Christos H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48:498–532, June 1994.
- [Per09] V. Perchet. Calibration and internal no-regret with random signals. In *Proceedings of the 20th international conference on Algorithmic learning theory*, pages 68–82. Springer-Verlag, 2009.
- [RST11] Alexander Rakhlin, Karthik Sridharan, and Ambuj Tewari. Online learning: Beyond regret. *Journal of Machine Learning Research - Proceedings Track*, 19:559–594, 2011.